AFRL-IF-RS-TR-2004-70
Final Technical Report
March 2004

PROBABILISTIC, DYNAMIC ANALYSIS OF PLANS

University of Massachusetts

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K442

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-70 has been reviewed and is approved for publication.

APPROVED: /s/
WAYNE A. BOSCO
Project Engineer

FOR THE DIRECTOR: /s/
JAMES A. COLLINS, Acting Chief
Information Technology Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	Mar 04	Final Apr 00 – Apr 03

4. TITLE AND SUBTITLE

PROBABILISTIC, DYNAMIC ANALYSIS OF PLANS

5. FUNDING NUMBERS

C - F30602-00-1-0529
PE - 64760E
PR - ATEM
TA - P0
WU - 12

6. AUTHOR(S)

Paul R. Cohen

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Massachusetts
Computer Science Dept
140 Governor's Drive
Amherst, MA 01003

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

DARPA
3701 North Fairfax Drive
Arlington, VA 22203-1714

AFRL/IFTB
525 Brooks Road
Rome, NY 13441-4505

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

AFRL-IF-RS-TR-2004-70

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Wayne Bosco, IFTB, 315-330-3578, wayne.bosco@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 Words)*

The University of Massachusetts developed three technologies under the Active Templates program. The first was Capture the Flag, a wargaming environment that includes a simulator of battalion-and-higher level land warfare. Capture the Flag was initiated by the Active Templates program manager LTC Doug Dyer, but funded primarily by other programs, so it is not reported in any detail here. The second project involved training rats to carry audiovisual backpacks under computer and human control for surveillance and search and rescue missions. The third effort concerned an hypothesis of LTC Dyer's that the semantics of the kinds one sees in the Active Templates program could be developed "bottom up" by the people who use them, rather than "top down" by trained ontologists.

14. SUBJECT TERMS

Active templates, semantics, animal behavior, search and rescue, surveillance, control of biological agents

15. NUMBER OF PAGES

30

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Table of Contents

Table of Figures

Acknowledgments

More than any DARPA project we have worked on recently, our Active Templates effort was guided by our program manager Col. Doug Dyer. The project had three major phases, each initiated or developed with Col. Dyer, and it broke new ground because Col. Dyer encouraged us to take research risks.

We are deeply grateful for the wise and steady guiding hand of LTG Charles Otstott during the first phase of the project. General Otstott informed every aspect of the *Capture the Flag* system, from the simplest facts about the size of an M1 tank to the subtlest details of our agent programming language. More than that, General Otstott gave us a strong sense of purpose and encouraged us to think of Capture the Flag as really useful technology.

The Capture the Flag phase of the project was not large and is not dealt with in this report, as funding for the project came primarily from other DARPA projects. We mention it here primarily to acknowledge Col. Dyer, who first suggested that since a computer had beaten the world champion at chess, it was time for a computer to win a war game; and Gen. Otstott, who showed us the way not to this goal – because war games are computationally much harder than chess – but close to it.

The second phase of the project, called *Packrats*, implemented the idea that rats could be trained to carry audiovisual backpacks under human and computer control and could serve in surveillance and search-and-rescue missions. The leader of this phase was Professor David Palmer, a psychologist from Smith College. In Prof. Palmer we found a unique combination of animal expert, hardware hacker, and philosopher. We did indeed train rats to carry backpacks, and we learned a great deal about the nature of cognition and behavior in the process. Much of the text of our report on the Packrats project was written by Prof. Palmer. Reading it, one gets an unusually clear sense of how the work proceeded week-by-week, a rare look at the daily ups and downs of research.

The third and final phase of the project was a lengthy dialog between Col. Dyer and me on *bottom-up semantics*, an idea Dyer advocated and I rejected initially. Preliminary experiments seemed to support the idea of bottom-up semantics, and I have since come to see it as one of the really important ideas in the development of knowledge systems.

Let me also acknowledge Clayton Morrison, who helped manage the Packrats project; David Westbrook and Gary King, the architects of Capture the Flag; and Andrew Hannon and Michi Oshima, who helped with the programming on Capture the Flag and Packrats.

Thanks to all for a deeply rewarding project.

December, 2003

Packrats

The Packrats Project trained rats to carry video and audio backpacks under AI and human control for surveillance and search-and-rescue operations. Rats are more capable in difficult physical environments than robots, but they are harder to control because they have their own goals and behaviors. The AI problem therefore is to plan the rat's activities and control the rat (who does not always respond as directed), re-planning as necessary, in real time, given limited sensor data. We were not able to achieve this goal, partly because we ran out of funding, partly because rats are less easy to control than we had hoped, partly because our technology was insufficient. What we learned, however, gives us confidence that projects like Packrats might succeed, and prove very useful, in future.

The original Packrats team, Prof. David Palmer, Prof. Paul Cohen, Prof. Carole Beal, and Dr. Clay Morrison set up the parameters of the study: Rats would be trained to respond to two tones, high and low. Low means "keep doing what you are doing," high means "begin a search behavior." Search behaviors would result in the rats changing orientation, and when the rat is oriented properly the tone would change downward. In this way the rat would be steered. We planned to broadcast the tones to the rat over radio frequencies. Rats hear very well so the tones could be all but inaudible to humans in the area and still be heard by rats through speakers mounted on their backpacks. In this way, rats could be controlled covertly. Indeed, because rats are thigmotactic (they scurry along walls) and don't like to be out in the open, they seemed ideal for surveillance, as well as for search and rescue in dark, inaccessible areas. The World Trade towers were destroyed early in the Packrats project and it was disappointing to us that the rats were not ready to serve.
Later in the project Prof. Palmer trained pigeons to fly to particular locations, anticipating that one day they would carry video cameras. Pigeons can see ultraviolet light, unlike humans, so in theory they can be trained to fly to illuminated target locations.

Early training

Two research assistants, Peicha Chang and Scott Howard, started work early in the summer of 2001 under the direction of Prof. Palmer. The plan was to start training pigeons and rats simultaneously, though along very different lines. The rat project received the higher priority, as Prof. Palmer deemed the chances of a successful demonstration by summer's end to be more realistic.

Rats are identified by the number of dye spots on their backs. Zero and one spot were "tone rats," and two and three are "carry rats." Zero and one were trained to eat from two feeders at opposite sides of a Skinner Box. For the first week they were trained in a

high/low tone discrimination in the box. Food was available contingent upon approaching and investigating the appropriate feeder, randomly chosen from trial to trial. The high tone was presented whenever the head of the rat faced the appropriate feeder; otherwise the low tone was presented. At the end of one week of training (6 and 5 days respectively for zero and one), both rats showed good control by the tone. That is, both rats would turn within five seconds upon hearing the low tone and proceed when hearing the high tone on nearly all trials. On about half of the trials turning was immediate. The cramped quarters of the Skinner Box facilitated acquisition, because reinforcement was never delayed more than a second or so from the onset of approach behavior during the high tone. But for the same reason, there were confounding cues; for example, the rat's position in the chamber at the onset of the high tone was fairly consistent. So on Monday, June 11, Prof. Palmer's assistants began training in a four-foot runway with a feeder at each end. The procedure was identical, as if the runway were a stretched-out Skinner Box. The first day was devoted to adaptation to the new apparatus and feeders. (The runway feeders, a different brand from the box feeders, make a loud noise when operated. This is helpful after adaptation but hurtful before.) After three days both rats showed good control by tone cues on about half of the trials, depending on the location of the rat at the outset of the trial, but competing behavior interferes for five to ten seconds on other trials.

On June 13, we began a time-out procedure for One-spot in the last third of the session: If he persisted in running in the wrong direction for a second or two, subjectively measured, we turned off the lights for a few seconds and began the trial anew. This appeared to work. By the end of the session the rat would falter in its erroneous course, and on several trials it reversed itself in mid-course. Once the rats behave properly in the straight maze, the plan is to move them to a kind of radial-arm maze (central chamber with many blind alleys, only one of which leads to food) rather than city blocks, as we think that contingency will provide especially good tone control.

The "carry-rats" were given several days of clicker training, in which a "cricket" noisemaker was paired with food in a one-feeder Skinner box. Then, using the clicker as a conditioned reinforcer, each rat was trained to approach and pick up a barbell made out of 12-gauge wire and plaster "bells." One rat chewed on the plaster, exposing the wire, and bloodied its mouth, so we bought a plastic barbell of perfect size at a pet supply place. Two-spot quickly learned to pick it up, and over the past few days we moved it out of the Skinner box and progressively to a 20-inch runway and a 48-inch runway. Two-spot learned quickly to run back and forth carrying the barbell. By mid-June, Scott began adding weight to the bells by adding BBs. Our goal is to get the rats adapted to carrying a three-ounce weight. Scott also painted the ends white and black so that the rat could be trained to approach the barbell from one side only (so that the camera will point ahead and not behind the rat!). Three-spot lags a little behind Two-spot. We planned to begin outfitting the two rats with a harness, as an alternative way of carrying the camera, but as of mid-June the animal-use protocol was not approved.

By mid-June we met with the person who was to build the electronics package the rats would carry. The package will be somewhat bigger than we expected. It appears that the first incarnation of the device might be easier for the rat to carry on its back than in its mouth.

It was apparent by mid-June that the hot/cold signals exerted good control under some conditions but not under others. If the rat were at the end of the runway at the onset of a trial, its behavior was well-controlled by the tone. If it were facing down the runway, we would present the high (hot) tone, and it would trot down to the other end and get fed. If it were facing some other direction, we would present the low (cold) tone, and it would quickly turn and head down the runway. So that was perfect. However, if at the onset of a trial the rat were trotting down the runway in the wrong direction, the presentation of the low tone had no immediate effect. The rat would continue running until it reached the end of the runway (the wrong end) before turning around, as if its momentum were a controlling variable. Also, one end of the runway evoked a lot of sniffing. Sometimes the rat would continue to sniff for a few seconds after the onset of the tone. Without a time-out procedure there was no penalty for either of these deviations from the ideal, so we introduced a brief blackout and reset the trial.

This worked beautifully for one rat. The new procedure was quite disruptive to the other rat, however, and some of his earlier gains disappeared temporarily. However, after another day of training, both rats are performing very well. They stop and turn around, even in the middle of a run.

Scott modified the barbell so that it weighs 40 grams, about half of its target weight. He ran the two “carry rats” and both were unfazed by the change in weight, but Prof. Palmer has misgivings about their ability to carry much more for any distance.

One of the rats will accept a harness, the other puts up a fight and bites Scott when he tries to fit it.

After four weeks of training two of our rats successfully acquired a high/low tone discrimination in a runway, so that when the low tone came on, they would reverse direction immediately from any point in the runway, even if they were hurrying in the contrary direction. In the presence of the high tone, they would quickly advance. This kind of precise control was exactly what we were looking for. But a runway is a highly specific environment, and while we might expect our control by the tone to generalize to something like a heating duct, we would not expect control to be maintained in a more open environment. So we began generalization training.

Prof. Palmer regards generalization training as the most formidable hurdle of the entire enterprise. Reinforcement strengthens the control of behavior by a particular setting and by similar settings. When the setting changes considerably, the target behavior may be evoked only weakly. Moreover, the probability of a particular behavior depends not only on its own history of reinforcement, but on the probability of competing behavior. We

would not expect a hungry rat to pass by a baloney sandwich on its left or an alluring female on its right just because a tone is on. But competing contingencies don't need to be conspicuous. Whenever a rat enters a novel environment, its repertoire of defensive and orienting reflexes is at full strength, and trained behavior might fall apart. (Rather like the novice actor who forgets his lines when he first walks out on stage in front of a real audience.) If time passes uneventfully, these reflexes habituate, and trained behavior might again emerge as the strongest behavior in the animal's repertoire. This problem of competing behavior in novel environments is more serious for us than for dog trainers, not because dogs are smarter, but because they are higher on the food chain. Natural selection has ensured that, relative to dogs, rats are worry-warts.

So there are two problems with novel environments: generalization decrement and competing behavior. The best that we can do to address these problems is to vary the training conditions, ideally to include those conditions under which the behavior is expected to occur in the future. (Other partial solutions are logically possible. For example, it might be helpful to surgically kill off the sense of smell, or to use electrical stimulation of the brain as a reinforcer, but these are not within the compass of our protocol, nor are they likely to be trouble-free procedures. See below for a report on a related project that uses direct brain stimulation to control rats' behavior.)

Our first generalization task was the radial-arm maze, a Plexiglas contraption about 4 feet in diameter, consisting of a central chamber with 8 arms extending from it. The rat was put in the central chamber, and food would be delivered when it reached the end of the target arm, randomly chosen from trial to trial. The tone continued to be reliably correlated with optimal behavior; that is, it was low except when the rat was oriented toward the correct arm. Initially, control by the tone was weak. Unlike the runway, in which mechanical feeders were remotely operated, the maze required hand-feeding, specifically, the dropping of a couple of pellets through a hole at the end of the target arm. This was awkward and tiresome to the experimenter, but more importantly, it imposed a short delay to reinforcement as well as a new stimulus startling to the rat, namely, the experimenter's arm looming over the rat's head on every trial. However, by the second day, performance improved markedly, and after a week of training, both rats were performing very well, making only incipient investigations of erroneous arms and moving ahead when the tone switched to high.

We were puzzled by some variability in performance, particularly when performance seemed to deteriorate from one day to the next rather than to improve. We still don't have an explanation for that variability, but continued training established the target relation in strength. One rat would systematically turn from one arm to the next, advancing only in the presence of the high tone. The second rat moved quickly in a wider arc. Sometimes this would make him pass over the target arm, causing a tone pattern of low-high-low. When this occurred, the rat would turn back toward the target arm. On occasions in which the rat found itself at the end of the diametrically wrong arm, the tone would change from low to high when the rat turned around; under these conditions the rat would

run straight out of the wrong arm, cross the central chamber without pausing, and enter the target arm. In short, the rats were now performing beautifully in both the runway and the radial-arm maze.

By the first week in July the rats had been moved to a T-maze: a four-foot runway terminating in a four-foot cross-alley with a feeder at each end of the alley. Performance is not yet optimal in the T-maze, but it is quite good. Our goal is for the rat to turn immediately when he makes a wrong choice, bringing on the low tone. On some of those trials in which the rat turns first in the wrong direction, he will go all the way to the end of the wrong arm before turning. This only takes a second or two, since the arm is only two feet long, but we require better control than that. To some extent, this problem of overshooting a choice point may be more of a problem in the lab than in a novel environment. The end of the alley is a familiar location that has been frequently correlated with food. Control by the tone has strong competition by prevailing stimuli. We faced the same problem in the runway and solved it by the use of time-outs. We are using the same strategy in the T-maze.

We have found that the absolute frequency of the tones is not critical, that the rats appear to respond to tone differences (high/low). We have only investigated this in passing, but it augurs well for switching control to ultrasonic (to us) tone frequencies. (Rats are sensitive to frequencies up to around 50 kHz, while humans can hear little above 15 kHz.)

By early July, 2001 the other two rats, the "barbell rats," are catching up with the tone rats. You may recall that we were using these two rats to explore the possibility of training rats to pick up and carry in their mouths a 3-ounce weight in the shape of a barbell. In the service of this goal, these rats were clicker-trained (that is, they were given lots of click-food pairings to establish the click as a conditioned reinforcer) and were fed by hand rather than by mechanical feeders. (These procedural differences might prove to be quite important in other respects.) We found that rats could readily be trained to pick up and carry a barbell, and they continued to do so as we added weight. Moreover, by painting the ends different colors, we were able to get the rats to pick up the barbell from the same side reliably.

The barbell rats had by July begun tone training as well, in both the runway and radial-arm maze. A history of clicker-training and hand feeding proved ideal for the radial-arm maze, since there is no automatic conditioned reinforcer at the end of "correct" performance, and since feeding must be done by hand. One of the two rats reached an excellent level of control: He makes an incipient movement toward each arm in turn and scuttles ahead when the high tone comes on. The remaining rat reached optimal performance in the runway and started radial-arm maze training on July 6.

By this point the project is likely to soon outgrow its 4' by 4' wooden maze. Unfortunately, regulations prevent us from letting the rats touch the floor of the animal

labs at Smith College. The College has a very high rating for its animal facility and guarantees all sorts of welfare for its animals, and apparently this welfare can be jeopardized by contact with the floor, presumably because *we* have contact with the floor!

Training in Open Environments

It took only a month to get the rats to perform beautifully under tone control in mazes. But the real challenges lie ahead, in novel, open environments. The project is generating a lot of interest among our colleagues and even our children, one of whom asked, "Once the rat has accomplished its mission, how do you get it back?" This and other problems depend on the rats' backpacks, which carry not only a video camera but also a radio receiver of the sounds we'll use to control the rats. Unfortunately, by midsummer, 2001, we had made little progress on the backpack. The technician in charge was swept up in preparations for her wedding and then got into a bad car crash, which she survived uninjured.

By mid-July, 2001, we had obtained permission to allow the rats to roam freely (albeit on construction paper) within the animal quarters. The idea of taking them to DARPA for a demonstration was nixed, however. Once they leave the animal quarters, they leave for good, lest they bring back disease or a taste for mentalistic psychology to other rats in the colony.

As to getting the rats back after a mission, our best guess is that they will operate within range of a homing signal. They could be trained to locate it; otherwise, we would have to guide them back. Rats can follow the scent of another rat, and presumably their own, but Professor Palmer doubts this would be an important source of control.

By July 30, 2001, the rats had completed eight weeks of training. Our animal use protocol for harness use was finally approved, so we ran the rats with an "empty" backpack: a Velcro band held on the rats' shoulders with rubber O-rings.

The use of the harness has led to some unpleasantness. Scott, Peicha, and Prof. Palmer all have been bitten at least once. But more importantly, one rat is out of commission for the rest of the summer. In one of our battles to install the harness his leg was broken, or at least badly sprained. Two of the remaining rats are now fairly docile when the harness is applied, though we have scars to remind us of their former opinion of the procedure. The remaining rat, objected to both having the harness put on and wearing it. We are the bosses, and we are usually able to get him to wear it, but we aren't looking forward to strapping on a camera.

Regulations in the animal facility preclude training the animals in novel settings or in hallways. As recipients of federal funding we must follow the guidelines, but they are in direct conflict with the scientific goals of the project, which are to train rats to follow our commands in novel environments. The success of generalization training is the most

important question in the project, but the regulations prevent us from getting a convincing answer.

By the end of July one rat had mastered a “city block” maze in which it had to negotiate a 48” by 48” maze made of arbitrarily placed blocks. Even though the rat performed well on most trials, there was some variability from day to day. Some days control by the tone was essentially perfect, and we were able to guide the rat around islands at our whim, but on others the rat would engage in some off-task behavior. In particular, the smell of food was a disturbance variable. The rat tended to spend time sniffing through the mesh of the maze at spots where it had previously been fed, and because we were hand-feeding him, the smell of food was strong at all times.

Meanwhile, Scott was running his rats on the floor of one of the experimental rooms. (We laid down paper to avoid regulatory objections.) The rat navigated an open environment dotted with clear plastic tubs in order to get to one of four feeders, randomly chosen. Performance improved over sessions and leveled off with good control on about 80% of the trials.

Around July 25, we moved to a large modular floor maze for all rats. When it became apparent that we were going to be unable to train them in a natural environment, we decided to make the largest configuration of passages that was reasonable in the experimental space available to us. Tone control is still somewhat variable, depending on how complicated the apparatus is, among other variables. We are puzzled by the variability.

At the end of July we decided to take Three-spot out of the tone-training program, partly because we have only one maze now and overlapping demand for it and partly because the rats are all telling us the same thing, and we think we can get more data by changing conditions. We are going to train Three-spot to approach the human voice. We have a room set up with four speakers and four feeders. The rat will be fed when it approaches the speaker that is on at any time. (We are using the tape of a conversation between B. F. Skinner and E. O. Wilson, recorded in 1987. We want our rat to be verbally sophisticated.)

At this point we also have a prototype device for picking up broadcast tones. It has room for a camera and a microphone too. Unfortunately, it is too big for our rats and has features that are beyond our needs.

Summary of the first ten weeks

We were unable to design a backpack to deliver tones to the rats; the one we commissioned was too cumbersome. We have still learned much about the feasibility of the project. We have shown that rats can be trained to lug barbells around and to wear a light harness and to be guided by a tone. We hope to show that they can be trained to

approach voices, at least as mediated by speakers. All this is good, but we have also bumped into some important limitations:

We have found one reliable thing with our rats that raises concern for the overall success of the project. Whenever we change conditions, tone control deteriorates dramatically. In the first session in a novel apparatus we always find that the rat does a lot of exploring and sniffing on its own, regardless of the tone. The behavior of the rat is under joint control of all concurrent contingencies. In a novel environment, all sorts of defensive and orienting reflexes are at high strength, and they compete with the target contingency. In a research project one would control for this by tightly controlling extraneous variables. One puts the animal in a sound-proof, enclosed chamber, with white noise to mask incidental noises. The animal is allowed to adapt to the apparatus for a session or two before the experiment is begun. (Skinner felt that most of the maze research that preceded his own work was worthless because too many variables were uncontrolled.)

But we have made a point of training the animal under highly variable conditions. Our apparatus is open to the "sky" of the animal quarters; there is full illumination and a lot of ambient noise; the rats can smell food and the traces of other rats; pigeons are coming and going into their Skinner boxes and drumming on response levers on fixed-ratio schedules, cooing and chortling as they do so. The current maze is in a large public room with people coming and going. It couldn't be worse from the point of view of experimental control. However, we are trying to introduce as much verisimilitude into our training conditions as possible so that generalization will be enhanced. And we have been heartened to find that the rats usually do well after they have adapted to the new conditions.

But novelty is an intrinsic part of any application, and novelty always disrupts our rats' performance, at least for a while. This is a topic that will require considerable discussion.

A movie of the Packrats project at this point in its development may be found at http://www.cs.umass.edu/~clayton/packrats1.mov.

Hardware problems

Our first commissioned backpack was a failure; it was too cumbersome, consisting of the inner board and components of a commercial headset, with the only modifications being a shortened connection to the speaker and a power supply added in. There are a lot of components that should be jettisoned (e.g., two power-indicator lights, volume dials, on/off switch, etc.).

We decided to switch course and construct a smaller rat-pack that consists only of a commercially-available, all-in-one transmitting video spy camera, and a power supply, and stick to broadcasting the guiding tone to the rat from "above" the maze apparatus (as is currently done in the experiments). This way, we could immediately collect data of the rat running around and being guided by tone, even though the tone source is not on the

rat's back. This will address our immediate questions of: whether the rat will tolerate a payload with components (as opposed to the current weighted "dummy-pack"), what kind of visual information we get, and whether having the camera on the back of the rat can be used for rat-perspective navigation.

We need to test as soon as possible how well we can guide a rat based on video transmitted by a camera harnessed to it. Thus, we worked to package our original all-in-one transmitter-camera with the smallest 9-volt power supply we could come up with (a standard 9-volt battery is too large and heavy for a rat).

Our first approach was to use three stacked 3-volt wafer lithium batteries. However, we found that, while the three batteries initially output 9 volts, they drop within 2 minutes to a lower voltage (around 5 volts), and the camera subsequently no longer transmits a viable signal. If you disconnect the camera, the voltage then slowly returns to around 9v. This indicates to us that the camera draws too much current from the batteries.

Clay Morrison also learned two important facts about Lithium batteries: they *will* explode if you get them too hot, and the magnesium oxide / lithium combination is really only toxic if you swallow the battery.

Our next approach is to connect the camera to a small 12-volt battery, with a resistor to bring the current down to 9-volts. Moving to this battery has two advantages: it should hold a steadier charge than the wafer batteries, and it is still significantly smaller and lighter than the standard 9-volt.

We outfitted power supplies (one 9-volt and one 12-volt battery) with the appropriate coax plugs to power and test the new components: the new audio-video transmitter, microphone, and new multi-lens camera. The good news is that they worked beautifully: we can transmit both audio and video clearly. And the range of the signal is good (we walked the unit down a long hallway around the corner from the rat lab with a strong signal until the very end). The new camera seems to transmit more clearly than the original all-in-one transmitter-camera (although this may be in part due to the higher quality of the new transmitter). Also, the sound is very clear -- we did some tests of talking while moving away from the unit, and performance was good. Once we remove all of the current RCA and coax plugs, they should be light and compact. We believe we should be able to power all three units using just a 12-volt battery if we can resist the current going to the transmitter down to its required 9-volts, but keep 12-volts powering the camera and microphone.

Further Packrats Research and Training

By September, 2001, one of the rats could be "steered" around the laboratory by tone control. Performance was not perfect, but on roughly half the trials we were able to start the rat in one of four rooms (off a common hallway), steer it into the hallway, and then into a randomly-chosen room, where it would be rewarded with a food pellet. The rat got

no information besides high and low tones, and these, when provided by a trained human handler, were sufficient to steer the rat as described. In these trials the rat wore a backpack with a video camera and broadcast unit, so for the first time we were able to see "rat's eye views" of the lab environment. (We realized right away that the imagery would require some correction, as it lurched up and down in a rather nauseating way as the rat ran along. Still, it was more than sufficient to recognize people in the environment. Success!

We briefed the work to Col. Dyer around Oct. 15, 2001, and later to Dr. Alan Rudolph at DARPA. Dyer directed us to use the remainder of our Active Templates money as we saw fit. Rudolph runs a project at DARPA on the "Control of Biological Systems." He's supporting Prof. John Chapin at SUNY Brooklyn on direct brain stimulation of rats -- both for motor control and reinforcement. Apparently they are doing well.

Around early November, 2001, we brought online the first version of "RatSim" a simulation of rats in a maze. For this we used the Tapir/HAC architecture that underlies Capture the Flag. We built RatSim as the first step toward having AI planners control rats: The idea is that they will first control simulated rats and, later, when we have an overhead camera in the Animal Facility, we'll try the AI with real rats.

The rats can be tracked by the Pioneer colored object tracking system that we use for our robotics work, provided we paint colored spots on their backs. RatSim simulated rats have the spots already.

We got to see Professor John Chapin's rats in mid-November, 2001, at the Southwest Research Institute in San Antonio. John Chapin's group uses two kinds of direct brain stimulation. Medial forebrain stimulation (MFB) is the reward signal, and two other electrodes plug into the region of the brain that senses the whiskers. The rats are trained to turn right and left when they feel artificial stimulation of the corresponding whiskers.

We saw the rats directed through indoor and outdoor courses. Two things were impressive: When put down in a new environment, the rats were immediately on task, no doubt because of the MFB. For example, I saw them put in a small pile of rubble, and they didn't explore a whole lot before they could be directed. The second impressive thing was the degree of spatial resolution in the control. They could be steered around a paint can, for example.

That said, there are many, many parallels between Chapin's project and ours, and I'm not convinced that their work is better, though the comparison really depends on the kind of tasks we have in mind.

One strong parallel is that the human controller has no easier time in their project than in ours. I had the opportunity to control one of their rats. They have three controls - left, right, and reward. It transpires that you have to reward the rats frequently, three or four times a second. Also, left and right are relative to the rat's frame of reference, and it was

very difficult for me to remember the orientation of the rat and send the right signal. Finally, after a few moments, the rat started to stagger alarmingly, and the grad student in charge raced over and took the controller from me. I have no idea what I did but I suspect it was too much MFB stimulation.

Prof. Palmer was absolutely right to suggest the hot/cold signal instead of a left-right signal. Interestingly, when it comes right down to it, Chapin's group is using a hot/cold signal, too. Hot is MFB stimulation, cold is the left or right signal. Personally, I don't know whether the additional discrimination (left or right vs. do something else) was worth the additional cognitive load on me, the controller.

I'm really interested in the codes we send to the rat. Hot/cold is one end of a spectrum, it's the most general, minimally discriminating signal one can send. What else could we do? The signal "vocabulary" has to mean something to the controller, and something to the rat; and we're after a coding system that is easy for people to use, easy for rats to understand, unambiguous and robust against noise, produces as wide a range of behavior in the rat as possible, and is interpretable to an AI. Another point on the spectrum is "hot/left/right," but that symbol system is difficult for people to use because sending the correct signal depends on the rat's orientation. On the other hand, if an AI could resolve orientation and send the right signal --- if the human could indicate a location in Cartesian space and the computer could turn this into a location in egocentric space --- then perhaps the range of behaviors evoked by the hot/left/right signal is worth the extra effort.

Another kind of signal is the auditory illusion – the idea of using stereo speakers to create the impression of a point in space to which the rat should orient. I'd like to do some good research on the design of an interchange language, for going from humans, through AI, to rats, so we have a theoretical reason as well as an empirical one for preferring this approach.

Another point of comparison: Chapin's rats carry 50-gram backpacks and can climb trees and chain-link fences. The rats are no bigger than ours (though most are female). I assume MFB stimulation is producing pretty dedicated rats. Interestingly, they don't fight the backpack, and are said to enjoy being run, no doubt because they associate the backpack with pleasure.

Prof. Chapin's rats need training just as ours do, the difference seems to be that they are more on-task; I suspect that their learning rates are not significantly higher, though Chapin did say that the rats could be taught to climb a fence, run over a 10' beam, etc. in a day or two.

While their rats are on-task to a remarkable degree (e.g., they were steered to a cheese danish pastry on the floor, and then steered away, poor things!), I don't think they generalize any better than ours. In fact, the grad students suggested training does not lead

to generalization. For example, one student taught the rat to approach speakers on a tabletop, but when the rats were put on the floor it didn't generalize.

It is difficult to get a clear picture of Chapin's rats' capabilities. They ran impressive obstacle courses, but I suspect they had been trained. They walked directly across a grassy field, but a grad student walked along beside them. They ran around a rubble pile, and the students said they saw it for the first time yesterday, but when I asked the student to move the rat from point A to point B, he wasn't able to. There was an awful lot of post-hoc explanation of what the rat had just done, and I didn't see very good control except in the "set pieces" indoors, on familiar tasks.

There are some advantages and disadvantages of MFB. John Chapin characterizes it nicely: You get immediate control of the immediate action, but this supervenes on any longer-term behavior, and effectively destroys it. You get moment-by-moment control, but you can't say "go over there" and let the rat figure out how to do it. In this context, "pulling" the rat to an auditory illusion is much easier, perhaps, than "pushing" it, like a shopping cart with a bum wheel. Chapin's group may not have the first option, because their control is moment-by-moment.

Another way to say this is that MFB stimulation really wrecks the behaviors that make rats attractive; it makes them much more like robots. We want rats because "native rat" and "trained rat" constitute two complete, coherent, capable, interacting levels of control. A rat under MFB stimulation isn't like this. In fact, John said very clearly that there are times when they turn off the stimulation and let the rats go back into "native mode," because sometimes that's what you want.

Continuing this theme, I think the great strength of our approach is the expertise we bring in the areas of training, AI control, and human interfaces. We can build a real four-level control system, it's hard to do with MFB stimulation. The attendees really liked the rat simulator and saw immediately the advantage of a human clicking on a location and letting the AI steer the rat.

John's team has a very small GPS system about to be deployed on the rat. He uses the same camera and transmitter as we do. We agreed that it would be very helpful to visit each other's labs. They like our stuff on controlling the rat by AI, pulling the rat toward auditory illusions, biasing rats to sweep wide areas, the dumbbell deployment, training to approach voices, and the generality of the hot/cold signal.

We got a lot of feedback, but I think much of it was confused. Many criticisms from military users were of the form, "I need X, you can't provide it." Rats cannot march at 4mph with Marines, they are tiny little animals and shouldn't be considered for large-scale operations. They cannot fly. Criticizing the work for reasons like this is pointless. On the other hand, we had champions as well. Dave Burdick, from the Naval Air Warfare Center, said "I can think of a dozen apps right now for the dumbbell mode of deployment." and "Can I train a dog to carry a rat?" Other ideas:

- Can rats find minefields. Can a rat drop a marker on a mine. Compare rat performance with dogs (Apparently there are multiple facilities for this kind of testing.)
- Demonstration of rat teams, a cooperative task
- Airdrop a habitat for a rat in a location, release the rat, surveil the area, do it every day. Drop the box from the air, leave the rat in place as ears and eyes.
- Building recon. Send in multiple rats with bugs. Let it be stochastic, or perhaps use communication repeaters in waves. Can we show that trip wires and other security can be avoided, etc? Get through security?
- Put an animal in a building and demonstrate to what extent it can cover the whole building, then recall the rat.
- Everyone agreed that rats have a role as first responders in disasters. The problem is that this stuff is funded by FEMA, not DoD, and no FEMA personnel were at the meeting.
- There were lots of UXO (unexploded ordnance) people in the room. They were here largely because of the honeybee sentinel work (which is pretty impressive --- the bees can detect and respond to parts per trillion). They want to know whether rats can smell esters from mines, etc.
- Someone pointed out that full video is very expensive in terms of power and is information glut anyway. The consensus seemed to be that snapshots would work just as well in many applications. Someone suggested flying something over and uploading imagery; pigeons, for example.

By early November, 2001, it was becoming clear that exploratory behavior in the rats was a problem. While direct brain stimulation seems to keep the rats focused and on task, our rats continue to do the things that rats do in new environments – sniff around and explore, look for food, and stay out of any situation that might get them in trouble with predators. Prof. Palmer reported, “Over the last two weeks I have talked with a bunch of professional colleagues, and I am persuaded that we can get better control than we currently have. On the other hand, I’ve learned that even dogs have the problem of competing exploratory behavior. One of the leading dog trainers in the business says that her dogs always do a lot of investigative sniffing when she introduces them to novel environments.”

The plan at this point (early November) is to try to solve two problems at once: The rats need more training, which is costly when humans are the trainers, and we want an AI system to control the rats. The idea, then, is to set up an overhead camera in the lab and get vision algorithms to track the rats, and planning algorithms to pick tasks for the rats and reinforcement signals. In other words, have a computer train the rats.

Around November 10, the focus of the AI contingent of Packrats was building a good rat simulator. The first version behaved to a first approximation like rats, but Prof. Palmer had additional suggestions: Once a rat starts off toward a “goal” such as a corner or doorway, it gets more “excited” about reaching that goal as it gets closer; as it gets more

excited, it is less likely to alter its behavior according to our signal. So rats are most receptive to guidance before they have fixed on some goal they are trying to reach, and get harder to control as they get closer to reaching their goal (after which time they “reset”).

Prof. Palmer wants the controller in the simulator work to control the rats full-time for training. He thinks that this would probably improve our current training results because there would be greater speed, consistency and accuracy in responding to the rat’s behavior. Also, we don’t know of anyone having used such a system to automatically condition animals.

The plan is to bring one of our cameras, a framegrabber, and the object tracking system (OTS) software to the Animal Facility at Smith and set it up in the lab to see what will be involved in tracking the rat while it is moving through a maze. The idea is to place two colors on the rat (blue and red), and use the Pioneer Object Tracking System (OTS) to track the rat -- the two colors are used to determine the orientation of the rat. If this works, we will then make OTS communicate with the controller, and have the controller be responsible for emitting the appropriate tones for conditioning.

Some additional components to this system that we have discussed: (a) Have the controller also control the food hoppers. (b) Perhaps also implement scoring criteria so that the tracking system + control could also collect and process training performance.
The tracking will give OTS a good workout. We don’t know how it will perform with the speedy rats. It might be necessary to purchase a higher-resolution framegrabber and other equipment.

Meanwhile work on the backpack proceeds with a new technical person.

During October we had made contact with a person who trains ferrets. Ferrets are predators (actually, they are vicious killers) so ought to be less skittish than rats. We will evaluate ferrets as alternative animals. We designed a hood for the ferrets which will eventually carry the camera. The rats at Smith continue to be conditioned and handled to keep their performance up. We are eager to get the autonomous training system up and running and see how it affects their performance.

Delays

We did not realize in November, 2001, that it would take months for the computer equipment and programs to work correctly. It was very difficult to set up an overhead camera, visual tracking software, and internal (planner) representations of the maze environment in which the rats were trained; as well as a planner to train the rats. By April, 2002, we had purchased and installed a dedicated machine (appropriately named Skinner), fashioned a power supply for the camera, and successfully grabbed clear, focused, color images from it. The plan was to get the new framegrabber to interface with our object tracking code, load up the new simulator/path planner onto the laptop

dedicated to running lisp and make sure it interfaces with Skinner, get the tracking system up and running with the new framegrabber, and test the map making code that builds internal representations of the maze given overhead imagery of it.

By the end of April, 2002, Skinner was generating tones and activating hoppers, the two things we needed to do in the physical world to train the rats. However, there were further delays and by the end of May, 2002, we still did not have an automated rat-training system.

It was July, 2002, before we had logging facilities and a version of the automatic training system. The logging facilities allow for logging of rat position and facing, tone generation, and hopper activation. The logging works for both automatic training and manual training (responding to control by an experimenter, including tone generation and hopper operation). The sound generation was re-implemented using a new sound library on the Linux side. The automatic training system still has some small bugs and potential improvements, but works in general.

At this juncture we decided to collect some data of rats moving around naturally and under tone control to test an algorithm and a hypothesis. The algorithm was called Voting Experts (VE) and had shown excellent performance on the segmentation task. Segmentation means cutting a time series into coherent "episodes" and we were eager to see whether it would segment time series of data from the rats into coherent "rat episodes." This was an exciting prospect because for decades psychologists have looked to a "natural" way to divide up animal behavior into episodes, and have always had to rely on their best guesses about the boundaries of these episodes. Now we had an algorithm that might do the job, it seemed worth testing the hypothesis that it could, especially as we were still waiting for the automated training system to work properly.

By late August, 2003, we had some results: Voting Experts returned behavior categories, but 95% of them are unique. We wanted to find *common* behaviors: if we hope to ever reinforce any of those behaviors they have to recur frequently.

Unfortunately, at this point the project ran out of steam. The funding was nearly gone, the automatic training system did not work well enough to provide consistent training for the animals, and some of the key people left the project.

Packrats: Conclusions

Although the project ultimately did not produce a search-and-rescue rat, it did produce a great deal of data and new knowledge about the prospects for controlled animals. Given the difficulties of working within stringent animal care regulations, having to build all our own hardware, and the small budget, the results are quite promising:

- Rats can be “steered” by tone control. After a couple of months of training, one can steer a rat from one room to another in a familiar laboratory environment using only tone control.
- Rats can be trained to seek out human voices in mazes. This result encourages us to think that rats may one day serve in search-and-rescue operations.
- Rats are a prey species and so are subject to a basic asymmetry: When a predator loses an encounter with prey, it goes hungry; when the prey loses, it loses its life. This asymmetry has profound effects on behavior: Rats are extremely cautious and do not adapt very quickly to novel environments. Survival behaviors dominate in novel environments, and we lose control of the animals until they settle down.
- Ferrets are predators, but they, too, are difficult to control, for symmetric reasons: When ferrets are put in new environments, they explore boldly, and this behavior overrides our control.
- Medial forebrain stimulation (as practiced by John Chapin) overcomes some of these control issues and raises others. Chapin’s rats are tightly-controlled – they will ignore a cheese danish on the lab floor when directed away from it – but they are essentially animal robots, and must be directed moment-by-moment.
- It is feasible to train rats automatically with a computer that sees what the rat is doing through an overhead camera, but it isn’t easy. Hardware and software issues were difficult to resolve and it will take more resources than we had to make the project work perfectly.
- It is feasible to build a video backpack for rats, and sometimes to get the rats to carry it. Our rats would occasionally “go on strike” and stop responding to tone cues when they tired of the harness. The backpack weighed just 30 grams, but it was a lot for the rats. Interestingly, under medial forebrain stimulation, Prof. Chapin’s female rats carried 80 grams.

In conclusion, we started the project with the dream of a four-level control system: At the bottom level are the rat’s innate behaviors, their extraordinary ambulatory skills and senses. The next level comprises the rat’s trained behaviors – seeking human voices, carrying barbells, running and seeking and stopping under tone control. The third level of control resides in a computer, an intelligent system with quick reflexes – far quicker than humans, who often were a “step behind” the speedy rats. One computer could control several rats again raising the hope of search-and-rescue animals. Finally, overall control would reside in humans. This four-level architecture involves three distinct kinds of intelligence. It remains a distant goal, but a highly motivating one.

Convergent Semantics

Can a group of individuals come to agree on the meanings of symbols simply by using the symbols in messages, or is a more directive, top-down oriented effort required? Luc Steels show how robots can come to agree on the meanings of made-up words by playing "language games." To one robot, a word might denote the size of an object, to another; the word might denote the shape of the object. Steels shows that the robots can converge on the meanings of words by using them in situations where their denotations are clear, as when we point to an object and name it. It is well-known that simply pointing and naming is not sufficient, as the word we use might denote not only the object but also any of its features, any of the actions in which it is involved, any relationships between the object and others, and so on. However, one's intuition is that the word plus the scene plus the opportunity to ask clarifying questions collectively are sufficient to pin down the meanings of words. Language games are rules of dialog by which one agent clarifies the meaning of an ambiguous word. Steels reported that communities of robots converge on the meanings of words through language games.

A great variety of language games is possible, and one would expect some to do a better job than others at producing consensual semantics. Col. Doug Dyer postulates "market forces" that move agents toward consensus about the meanings of symbols. The question we explored together is this: "What is the minimum language game sufficient to drive agents toward consensus?"

The Scene

Agents communicate about something called the *scene*. Throughout this note, the scene is very simple, a number line divided into regions. The *range* is an interval [0…N] *sub-ranges* comprise a set of possibly overlapping contiguous elements [$i \geq 0$ … $j \leq N$]. The *vocabulary* is drawn from a set of discrete symbols **W**={A,B,C,…Z}.

At the beginning of an experiment, each of M agents establishes a *mapping* between symbols in **W** and sub-ranges; for example, one agent might establish this mapping:

```
((A 0 8) (B 9 17) (C 18 23)
(D 24 36) (E 37 39)(F 40 51)
(G 52 62) (H 63 81) (I 82 84)
(J 85 100))
```

Meaning that A denotes the sub-range [0…8], B the sub-range [9...17], and so on. Then the agents send messages to each other. The rules for sending messages, the contents of the messages, and the actions taken in response to the messages can all be varied experimentally.

A Measure of Semantic Coherence

The token A denotes a sub-range for each of M agents. If all M sub-ranges for A were identical, then we'd say the agents agreed perfectly on the meaning of A. In general, though, the sub-ranges will be different; for example, one agent might think A means [7...15] while another thinks A means [10...24]. The distribution of the lower and upper bounds of the sub-ranges for a word are a measure of the agreement among the agents about the meaning of the word. In the previous example, the distribution of lower bounds is {7, 10} while that of the upper bounds is {15,24}. If agents agree on the meaning of a word then the distribution of lower bounds should have a small variance, and so should the distribution of upper bounds. Variance is just the sum of squared deviations of elements in a sample from their mean. Let $SS(w)_l$ denote the sum of squares for the distribution of lower bounds for word w, and let $SS(w)_u$ denote the corresponding sum of squares for upper bounds. Then degree of disagreement among agents about the meaning of word w is

$$Disagreement_w = \sqrt{\frac{SS(w)_l + SS(w)_u}{2M}}$$

The degree of *semantic divergence*, Δ, is the average per-word disagreement, that is,

$$\Delta = \frac{\sum_w Disagreement_w}{\text{number of words}}$$

We can also define *semantic convergence*, χ, as

$$\chi = \frac{1}{log\Delta}$$

In the following experiments, agents send each other messages of the form "word, number". The message means, "In my language, this number is denoted by this word." For example, "K,27" means, "In my language, 27 is part of the sub-range denoted by the word K." Let **S** be the sending agent and **R** be the receiving agent. Let $l(w)$ denote the lower bound of the sub-range denoted by word w, and $u(w)$ be the upper bound of the sub-range. So if the word K denotes the sub-range [15,31] to me, then $l(K) = 15$ and $u(K) = 31$ for me.

S composes a message as follows: It selects a number n in the range $0...N$ randomly, then finds the sub-range that includes n and the word w which denotes this sub-range, yielding the message "*w, n*". A recipient **R** is selected at random from among the agents and the message is dispatched.

Experiment 1

In this experiment there were 50 agents. At the outset of an experiment, each agent divides the range 0...100 into ten sub-ranges. Repeatedly, a sender **S** is selected randomly from the agents and so is a receiver **R**, and a message is composed and sent as described above.

When **R** receives *"w, n"*, it applies the following rule: If *n>u(w)*, increase *u(w)* by one. If *n < l(w)*, then decrease *l(w)* by one. Otherwise do nothing. In other words, **R** stretches its definition of w a little to bring it closer to the definition of **S**, the sender, unless i is already within the sub-range that **R** denotes by *w*.

The Δ score (defined above) is calculated after each epoch of 100 messages. Figure 1 shows how the score increases over 500 epochs. Clearly, as more and more messages are sent between the agents, they agree more and more on the meanings of words. This is good. Moreover, the decrease in Δ is most rapid in earlier epochs, also good – much of the agreement between agents is reached in the earlier epochs. However, in the second panel of Figure 1 we see that the definitions of the words tended to overlap a great deal. This figure shows on the vertical axis a particular word, and on the horizontal axis a check mark for the upper and lower bounds for the word for each agent. The "spread" along a row thus represents both the variability of the bounds and the tightness of the definition of the corresponding word. One can see easily that the definition of the first word, for instance, overlaps that of several other words.

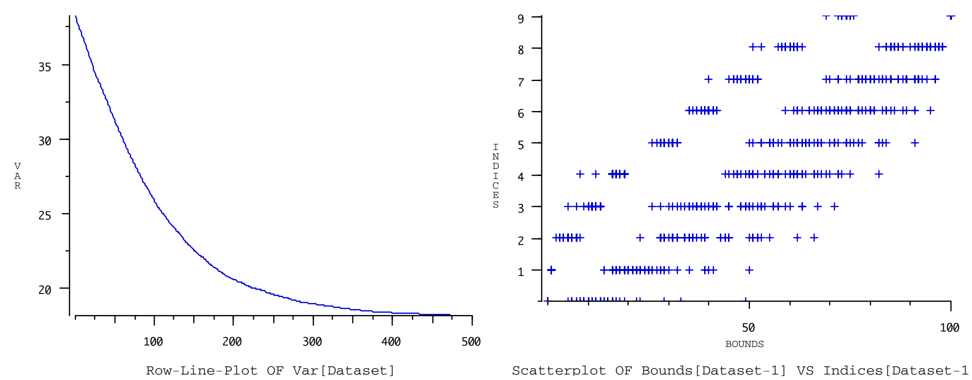

Figure 1: Results of Experiment 1.

The disagreement statistic Δ decreases in value as agents exchange more messages. However, the definitions of the ten words in the experiment overlap a lot.

In this experiment, 50,000 messages were sent between 50 agents, so each agent received on average 1000 messages. These messages referred to 10 sub-ranges or 20 bounds

(upper and lower). It seems to me that the convergence rate is quite slow, as it took roughly 1000 messages at each agent to get a degree of agreement on 20 things. Even after the disagreement rate Δ flattened out (near what one assumes is the asymptotic minimum) the word definitions overlap a lot.

Experiment 1 shows that simple rules for adjusting word definitions locally will produce semantic converge, but slowly, and the word definitions are not very precise. But it's a start, and we must now consider what is necessary to improve both the precision and the rate of converge.

Experiment 2

Let's suppose an agent remembers all the word, number pairs it receives. Then it could take the mean value of the numbers associated with a given word as a sort of "prototypical value" for the word. For instance, if the agent receives "D,14" followed by "D,18" it could calculate that the mean value associated with the word D is 16. However, central (mean) values are not the same as word definitions: A word definition has an upper and lower bound for the word, a sub-range of values that the word denotes. How can the agent decide on a sub-range? Suppose the agents agreed that a word definition should be a symmetric interval of some width around the word's mean value. Ideally the interval width would not be of fixed size but rather would depend on the messages that the agents pass to each other. Confidence intervals have this property. The confidence interval is based on a quantity called the *standard error*, defined as follows:

$$\hat{\sigma_{\bar{w}}} = \sqrt{\frac{s_w^2}{N_w}}$$

s^2_w is the variance of a sample of numbers associated with the word *w* and N_w is the size of the sample. Clearly, whenever the agent receives a message *"w, n"*, the value of s^2_w changes and N_w increases by one. In general, as the agent receives more messages about word *w* the standard error gets smaller. The definition of a word *w* for an agent may then be represented as a confidence interval

$$\bar{w} - k\hat{\sigma_{\bar{w}}}, \bar{w} + k\hat{\sigma_{\bar{w}}}$$

where the first term is the mean of the number associated with *w* in messages received by the agent. Note that *k* is the only free parameter in the confidence interval. It represents how "tight" the interval will be. Small values of *k* will produce small intervals. I used *k=2* in the following experiment.

Messages are generated as in Experiment 1, the only thing to change is what happens when an agent receives a message. Simply put, when the agent receives *"w, n"* it adds the number *n* to the sample of values it already has for the word *w*, recalculates the mean,

variance, and confidence interval for the sample, and sets the lower and upper bounds on its definition of *w* to be the values in the previous equation.

The results for this experiment are much better than those for the previous one. Figure 2 shows that the agents quickly converge on word definitions. The Δ statistic, which measures disagreement, decreases more quickly than in Experiment 1 and probably has a lower asymptotic value. It is gratifying to see that the word definitions do not overlap to anything like the degree they did in Experiment 1. Indeed, the words both span the range of discourse and overlap hardly at all. Said differently, few points in the range do not have a word to refer to them and few points are referred to by more than one word.

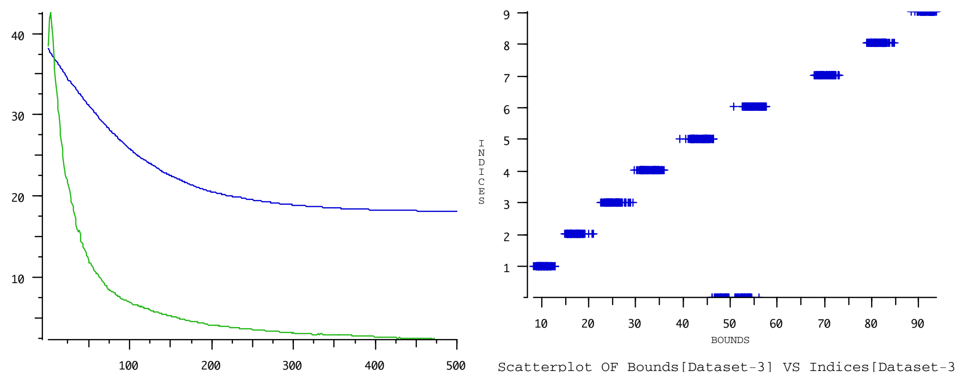

Figure 2: Results of Experiment 2.

The confidence interval method converges much more quickly than the method of Experiment 1 and settles at a lower disagreement level. Moreover, the definitions of the ten words overlap very little.

Still, the experiment involved only ten words, so I increased the number of words to 26 and re-ran the experiment with the same value of *k*. The results are shown in Figure 3. As expected, the convergence rate was slower than in the previous trials (the curve for these trials is the one in the middle). Remarkably, the word definitions appear to be as precise (i.e., not overlapping) as they were in the previous trials: Each word now has a smaller sub-range.

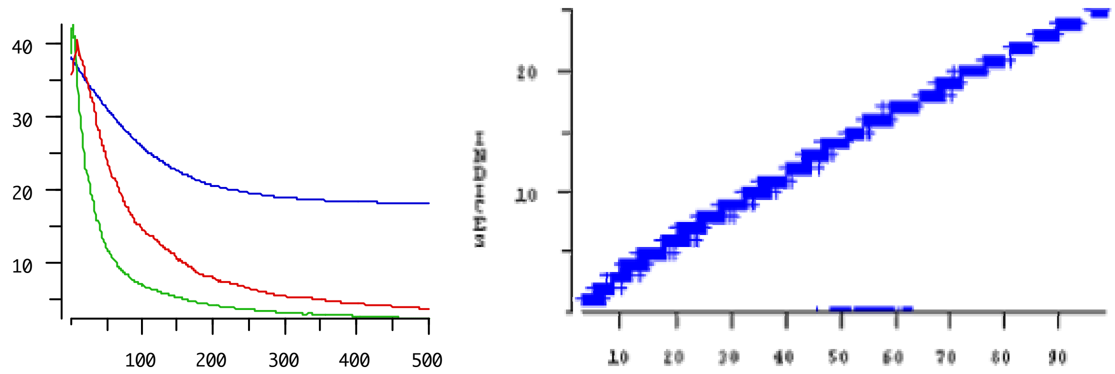

Figure 3: Results of Experiment 3

While the rate of convergence to low values of Δ is slower with 26 words, it is still faster than in Experiment 1, and the overlap between the word definitions remains low, despite that fact that the only experimental parameter to change was the number of words

This happy result appears to be a consequence of using the confidence interval as a word definition. The previous equation tells us that irrespective of *k* the size of *my* interval for word *w* depends on the variability of *other* agents' use of *w*, which depends in turn on the size of the interval for *w* for other agents. There's probably an interesting theorem here to the effect that the interval must converge to its correct value.

Discussion

Let's review progress: I defined a "scene" to be the range 1...N and a word denotation to be a sub-range. I asked whether agents could converge on a consistent set of denotations and defined a measure of disagreement, Δ. In Experiments 1 and 2, I varied only how the agents update their word definitions when they receive messages of the form "word, number". The method in Experiment 1 worked poorly: Although agents did converge on word meanings, the convergence was slow and the meanings were imprecise in the sense that several words referred to any given point in the range 1...N. The results from Experiment 2 were much more promising.

Experiment 2 shows that agents can converge on word meanings in an extremely simple "language game." All communicative acts are one-way and each involves a single message. Contrast this with Steels' language games in which one agent sends, another receives and then sends a reply that indicates whether it understands, then the first agent sends a clarifying response. My experiments aren't directly comparable with Steels', as his agents communicated about a more complicated scene, yet I wonder whether the clarification dialog in Steels' language games is strictly necessary.

- While my scene is as simple as it can be, a range of numbers, it should be easy to extend the methodology here to more complex scenes. Here are some obvious and immediate extensions:
- Multiple real-valued dimensions instead of just one; messages use words to denote regions of this multidimensional space.
- Unequal density of points in the space, so agents are more likely to send messages about some regions than others. In this case one would expect words definitions to correspond to the boundaries of dense regions. Perhaps no words at all would emerge to denote sparse regions.

In the current experiments, k was a parameter. If we assume that words denote regions in a multidimensional space it should be possible to learn k for each word, so some words are more tight or precise than others.

We are still a long way from solving the problem that motivated this work. Col. Doug Dyer's vision is to have people send each other simple, structured messages about everyday (military) things and to ensure quick agreement, via the passage of messages, on the semantics of the words in these messages. I've shown that agents can converge quickly on the meanings of messages about something very simple, the range 1...N. I need to formulate the problems Doug Dyer describes in terms like those in this note to test experimentally whether his vision of "semantics by example" is possible.

Further Analysis

Is the "meaning of a term" simply all possible values that term could take on? I.e., is the meaning of a variable its domain? When Col. Dyer talks about the advantages of structured data, he means that there are constraints on variable values, and these constraints are an important part of meaning (relationships between other variables also being important). Semantics-by-example provides information about the known part of the domain. When Jim Hendler says, "A little semantics goes a long way," he's describing the beneficial effect of knowing part of the meaning. Rule bases that are incomplete describe some of the relationships---and sometimes these are useful.
What does my simple scenario assume?

1) meaning is a denotational relationship between words and things in the environment (words denote sub-ranges);
2) a notion of "locality of meaning" (a word denotes consecutive values in a sub-range, not randomly-distributed values);
3) language is mediated by mental structures we call concepts (a sub-range is a concept, a discrete representation of part of a continuous world about which agents may speak);
4) a word, you already know part of its meaning, but you and I might not agree on all its meaning (when you receive E,15, you know what E denotes to you, and you learn part of what it denotes to me, but you don't know all of what E

denotes to me.)

5) each agent makes hard (as opposed to fuzzy or probabilistic or "graded") boundaries between word meanings (e.g., 13 is the upper limit of, say, D and 14 is the lower limit of E.)
6) weak semantic ambiguity in communication (when the sender says "D,7" the receiver knows that D refers to 7. The receiver doesn't know the entire meaning of D to the sender, but it does know that D denotes 7. If the sender said "D, 7 or 3 or 22 or 81" then the communication would be strongly semantically ambiguous;
7) the world to which words refer is a single, scalar dimension;

Assumptions 1 - 4 are fine by me; relaxing them makes the situation less not more realistic. Assumption 5 is wrong psychologically; human categories are graded. Assumptions 6 and 7 are problematic and are related to each other and to the structure of the scene.

Even in the trivial environment I've described, one can have richer concepts *and* semantic ambiguity *and* locality of meaning. *Structure* is what makes this happen. For example, there might be a concept of being near the middle of the range, or being near zero; or there might be a concept that represents a window of width two around a boundary. Let's consider the latter because it is interesting. Suppose I have three sub-ranges, 0 - 5, 6 - 10, and 11 - 21, denoted by A, B and C; and I also have the concept I just mentioned, for which I use the word D. To me, the regions 5,6 and 10,11 are denoted by D, because they are regions of width 2 around sub-range boundaries. Suppose that you have exactly the same *ontology* as I, that is, you also have words that denote sub-ranges and words that denote the boundaries of sub-ranges, even if we're not in complete agreement about what our words denote. So you know that a word from me might denote one kind of concept or another. Now suppose you receive D,5; what can you conclude about what D means to me?

You don't know whether D denotes a sub-range or a boundary region – you don't know which aspect(s) of the scene a word denotes, there's semantic ambiguity – but you can start to figure it out! If you later receive (from me) D,81 then by the principle of the locality of meaning you know D refers to a boundary region. Or, if you already know roughly what A means to me, then you can guess that D,5 denotes the boundary of region A. Or, if you've already guessed that I'm using D to denote boundary regions, then you can guess that 5 is the boundary of one of my words. And of course you know which word *you* use to refer to 5, perhaps it's A, so you can tighten up your hypothesis about what my word A means without ever receiving the word A from me.

How is all this possible? Structure! The principle of locality of meaning refers to the structure of the environment. The fact that we have different concepts --- sub-ranges and boundary regions --- reflects the structure of the environment. The fact that we can use a lot of knowledge to infer what the sender must have meant is due, again, to structure.

And, on the flip side, structure in the environment means words can be semantically ambiguous because they refer to this or that part of a scene.

But here's a theorem, a key theorem for Dyer's theory of semantics by example: While semantic ambiguity slows down the rate of convergence on word meanings, the structure from which it arises speeds up the rate of convergence, so structured environments lead to semantic coherence faster than less-structured ones.

I don't completely understand the relationship between structure, semantic ambiguity, and convergence on word meanings, but we have the same intuition that meanings are easy to induce when most of the scene is understood. Case in point: I say to Allegra "please put the cocktail shaker back in the cupboard." She doesn't know it's a cocktail shaker, but it's the only thing on the table that she *doesn't* know, so that must be it.